National Aeronautics and Space Administration

HIGH-END COMPUTING AT NASA

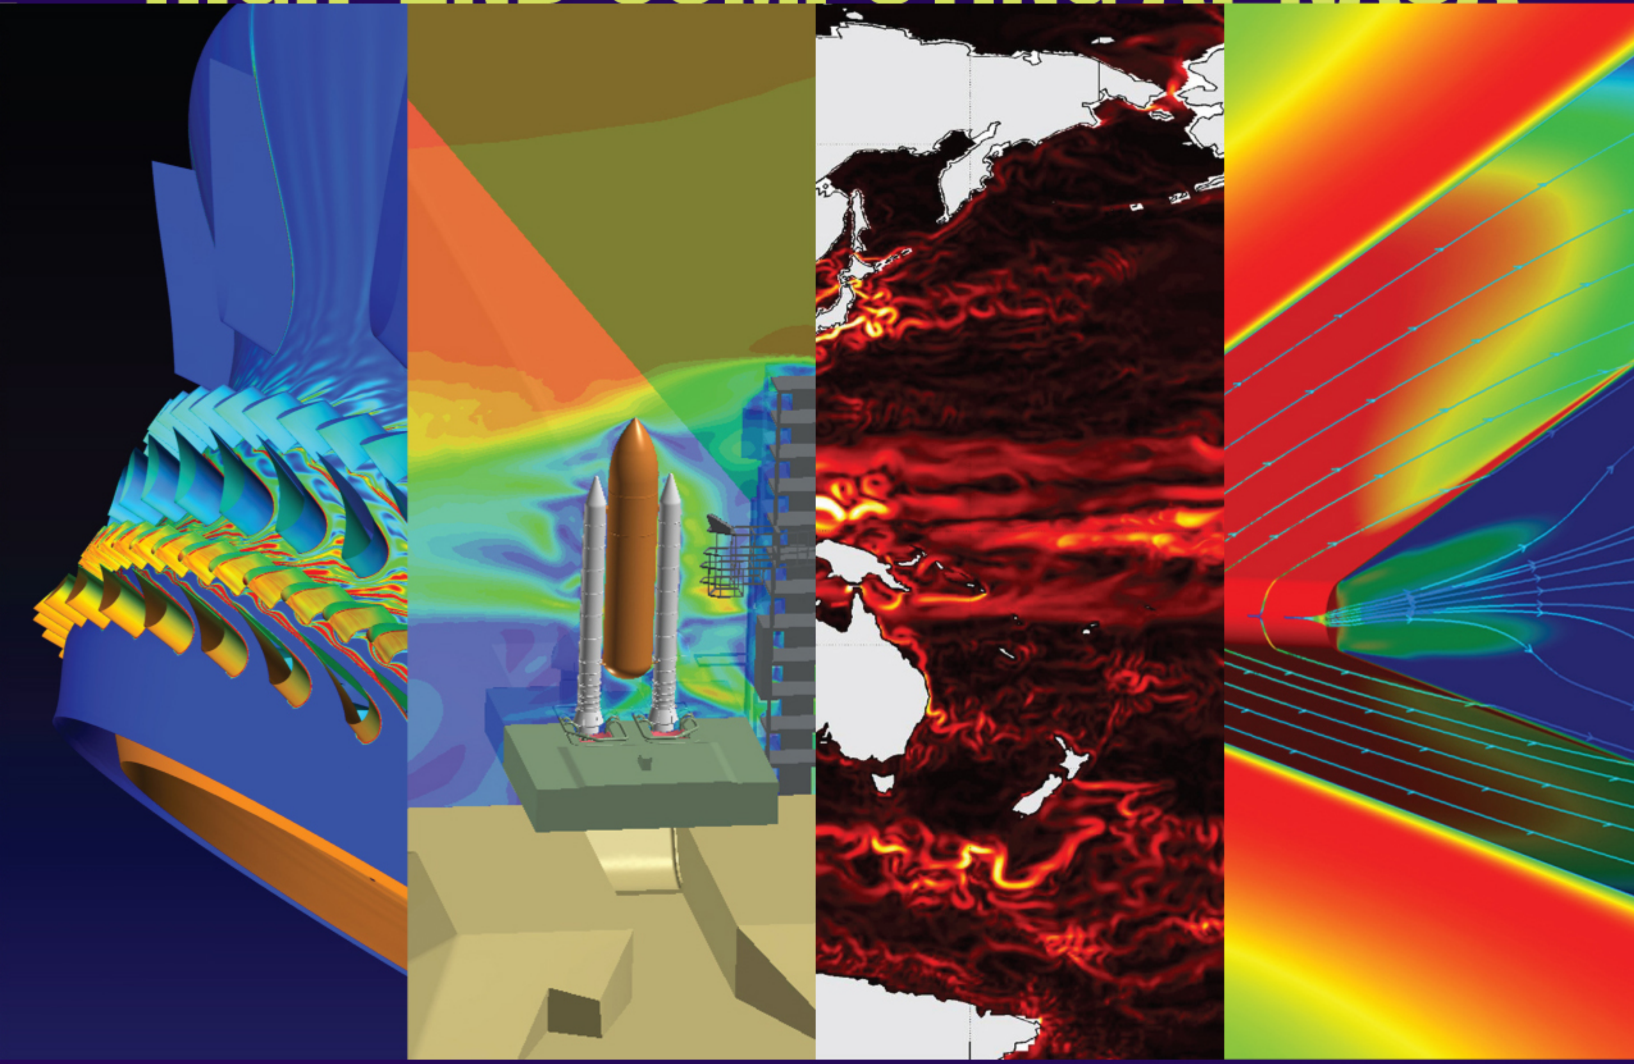

Highlights of science and engineering achievements on NASA's supercomputers

2007–2008

HIGH-END COMPUTING

VISION STATEMENT

For NASA's science, engineering, and technology community to consider high-end computing resources and services as essential tools that broadly enable rapid advances in knowledge and insight, and dramatically enhance mission achievements for the Agency.

Images shown on front cover (from left to right):

- Instantaneous view of the 5 blade-row high-pressure aircraft engine turbine coarse mesh simulation with the flow colored by vorticity to show velocity non-uniformities. The blade surfaces are colored by static pressure.
- Velocity magnitude color contours on a vertical plane passing through the centerlines of both Space Shuttle solid rocket motors.
- A global, coupled ocean-ice model simulates variables such as ocean current speed (shown here) and ice concentration.
- Sample wedge in an arc jet flow on the Orion crew exploration vehicle thermal protection system.

PROGRAM LETTER

June 1, 2009

NASA High-End Computing Community and Stakeholders:

We are delighted to present this report on NASA's High-End Computing (HEC) Program, covering the years 2007 and 2008. This publication captures significant science, engineering, and technical achievements from across NASA, enabled by the Agency's world-class high-end computing resources and services.

High-end computing serves an increasingly important role in NASA's missions—helping to safeguard our Space Shuttle fleet and astronauts, design next-generation space exploration vehicles, advance understanding of human impact on Earth's climate, expand knowledge of the origin and evolution of our universe, improve aerospace modeling, and much more. The HEC Program continues its commitment to maintaining a productive, service-oriented computing environment for all of these endeavors.

In this report, you will read about the technologies and services that make the HEC Program increasingly essential to NASA missions, and about the pioneering work that these resources have enabled. The report describes the Program's integrated environments that include premier computing systems, high-speed networks, a vast data management and archive capability, and an array of support services. During the past two years, the Program's facilities have achieved a 10-fold increase in computational capacity. Such technological accomplishments support principal investigators (PIs) from all four NASA mission directorates. In this report, 43 of these PIs relate their HEC-enabled science and engineering successes.

As NASA expands the frontiers of space exploration, scientific discovery, and aeronautics research, the HEC Program continues its dedication to providing computing resources and services to advance the Agency's missions. By delivering a reliable, service-driven computing environment that maximizes scientific discovery and engineering productivity, we will help ensure the success of America's space program for generations to come.

Dr. Tsengdar J. Lee
Scientific Computing Portfolio Manager
NASA Science Mission Directorate
NASA Headquarters

Dr. Rupak Biswas
High-End Computing Capability (HECC)
Project Manager
NASA Advanced Supercomputing Division
NASA Ames Research Center

Dr. W. Phillip Webster
NASA Center for Computational Sciences (NCCS)
Project Manager
Computational and Information Sciences
and Technology Office
NASA Goddard Space Flight Center

TABLE OF CONTENTS

EXECUTIVE SUMMARY

High-fidelity modeling and simulation, enabled by supercomputing, are increasingly important to NASA's mission "to pioneer the future in space exploration, scientific discovery, and aeronautics research." While scientific and engineering advancements used to rely primarily on theoretical studies and physical experiments, today computational modeling and simulation are equal partners in such achievements. As a result, the use of high-end computing is now integral to the Agency's work in all four mission directorates.

NASA's High-End Computing (HEC) Program provides high-level oversight and coordination of the Agency's two HEC projects: the High-End Computing Capability (HECC) Project implemented by the NASA Advanced Supercomputing (NAS) Division at Ames Research Center, supporting HEC users in the four Mission Directorates; and the NASA Center for Computational Sciences (NCCS), implemented by the Computational and Information Sciences and Technology Office (CISTO) at Goddard Space Flight Center, supporting HEC users in the Science Mission Directorate. Funded by the Strategic Capabilities Assets Program and the Science Mission Directorate, these projects facilitate hundreds of computational projects from across the Agency. The HEC Program's Board of Advisors represents the strategic interests of each mission directorate.

This report presents the Program's successes in dramatically enhancing its facilities, resources, and services to meet the mission directorates' escalating demands, and describes advancements that Agency scientists and engineers have made using HEC Program resources in 2007 and 2008.

The *Program Overview* section (page 3) describes the range of resources and support services provided at the two HEC Program facilities and how they benefit the Agency's high-end computing users. Integrated HEC services are essential to making the Agency's supercomputing resources operate at peak efficiency and to helping users be as productive as possible. Following are some of the major technical achievements over the last two years.

- Facilities: The NAS facility has undergone significant expansions and upgrades—almost doubling its computer floor space and power and cooling infrastructure to accommodate new supercomputers. NCCS completed the necessary support infrastructure upgrades for a major increase in computing power and a six-fold increase in storage.
- HEC Operations: NAS installed the Pleiades supercomputer, currently the third fastest in the world. NCCS increased their supercomputing capacity five-fold with installation of the Discover system.
- Data Management and Archive: NCCS is expanding scientific collaboration and advancement through development of a portal for intelligent data access and use.
- High-Speed Networking: A joint NAS-NCCS effort resulted in a 54-fold speedup in data transfer for a major weather modeling project.
- User Support: NAS supported engineers doing time-critical launch and reentry computations for Space Shuttle missions, and saved 2 million processor-hours by optimizing three important aerodynamics codes. NCCS supported U.S. and international scientific field campaigns, nationally important climate research, and spacecraft environment simulations for mission engineering efforts.
- Data Analysis and Visualization: In spring 2008, NAS debuted the hyperwall-2 visualization system, enabling unprecedented large-scale data analysis and concurrent visualization. NCCS installed an interactive, large-memory data analysis system that gives users direct access to NCCS' 1.2-petabyte global filesystem and data archive.

NASA's HEC users currently number over 1,500 and come from every major NASA center, as well as universities, industry, and other agencies. The *Science and Engineering Highlights* section of this report (page 19) features the accomplishments of over 40 computational projects, selected based on their impact to the Agency over the past two years. For example:

- Aeronautics Research: Engineers are modeling innovative space exploration vehicles, aircraft components, and combustion and propulsion technologies. A Boeing team assessed the aerodynamic characteristics of the X-51 hypersonic vehicle, scheduled to fly in fall 2009. This project extends previous work completed by NASA on the X-43 Program in pursuit of practical hypersonic flight.

- Exploration Systems: NASA is designing America's next-generation spacecraft to take humans back to the Moon and ultimately to Mars. Engineers at Marshall Space Flight Center used over 7 million processor-hours to simulate the J-2X engine, which will power the upper stages of the next-generation Ares I and Ares V vehicles.

- Science: Researchers are advancing models and analyzing observations to better understand Earth's system and its impact on humankind, our planet's relationship with the Sun, and the evolving solar system and universe. Astrophysicists from NASA Goddard made fundamental discoveries about black hole mergers that are essential to the success of the Laser Interferometer Space Antenna, the first instrument expected to directly measure gravitational radiation from space.

- Space Operations: Research and engineering activities support the Space Shuttle and International Space Station programs, as well as launch, space transportation, and space communications work in both human and robotic exploration programs. NASA Johnson engineers led a project to model the ascent aerodynamics and debris transport of the Space Shuttle to protect Agency assets and, most importantly, its people.

The report also looks at future directions in Agency use of HEC technologies over the next several years. The Program will support strategic directions set for NASA by the new administration to further climate change research and monitoring; mount a strong program of human and robotic space exploration; support safe flight of the Space Shuttle to complete assembly of the Internation Space Station; and continue NASA's commitment to aeronautics research. The HEC Program's continued dedication to providing the best performance, usability, productivity, and security of its resources and services will propel computational science and engineering advances and help assure success for NASA missions.

PROGRAM OVERVIEW

INTRODUCTION

Over the past several years, NASA's mission "to pioneer the future in space exploration, scientific discovery, and aeronautics research" has been greatly enhanced through the contributions of high-fidelity modeling and simulation, powered by the best available supercomputing resources. As computational analysis is now an equal partner with theoretical study and physical experiments, high-end computing (HEC) has become an integral part of NASA's efforts in every key mission area. Major advancements are increasingly emerging from computational modeling first, and later validated by experimental studies and observations, or explained by theoretical work.

NASA's HEC Program, now in its fourth year, provides high-level oversight and coordination of the Agency's two HEC projects: the High-End Computing Capability (HECC) Project, formed under the Strategic Capabilities Assets Program (SCAP), implemented by the NASA Advanced Supercomputing (NAS) Division at Ames Research Center, and supporting HEC users in the four Mission Directorates; and the NASA Center for Computational Sciences (NCCS), implemented by the Computational and Information Sciences and Technology Office (CISTO) at Goddard Space Flight Center, and supporting HEC users in the Science Mission Directorate (SMD).

This program-level coordination ensures that the HEC facilities managed by these two projects provide a comprehensive set of supercomputing resources and services addressing the requirements of NASA, its external collaborators, and the nation. Overall HEC resource allocations are made annually to the mission directorates by the Program, and each mission directorate sub-allocates and prioritizes its own computational projects. A HEC Board of Advisors represents the strategic interests of the mission directorates, SCAP, and the office of the chief information officer. The Program is managed by SMD and funded through SMD and SCAP, which receives funding from all mission directorates to serve as an Agency-wide resource.

Since the 2006 *High-End Computing at NASA* report, usage of the Agency's HEC resources has grown rapidly—to more than 119 million processor-hours per year. During that time, the Program has also increased its overall computing capacity more than 10-fold. This expanded capacity has allowed the HEC Program's more than 1,500 scientific and engineering users at NASA centers, government laboratories, universities, and corporations to tackle nearly 500 computational projects in diverse disciplines.

The impact of users' computational work on mission success has never been greater. For example, Space Shuttle Orbiter support engineers have been simulating the transition from laminar to turbulent flow caused by protrusions such as gap fillers or thermal blankets. Using NASA's Columbia supercomputer, these engineers can quickly compute and predict the heating effects of such protuberances on the shuttle's thermal tiles during reentry into Earth's atmosphere. This capability is vital for real-time shuttle safety assessments during a mission. Moreover, such simulations are being used by the Exploration Systems Mission Directorate to optimize the design of future spacecraft. Increases in available supercomputing resources have also enabled NASA scientists to address their key Agency role in the national climate research initiative.

Achievements such as these rely on the tools provided by a comprehensive computing infrastructure including: HEC operations; data management and archive; high-speed networking; user support and application optimization; and data analysis and visualization. The pages that follow describe these services—provided by the NAS and NCCS facilities—along with examples of their positive impact in solving NASA's scientific and engineering challenges.

FACILITIES

The NAS and NCCS facilities are home to the HEC Program's most valuable assets—its computing resources and the people who support those resources and their users. Together, these facilities provide NASA users with more than 800 teraflops of computational power from six systems, distributed across 40,000 square feet of computer room floor space. In addition to high-end computing systems, each facility houses auxiliary equipment required to run the systems including front-end systems, disk arrays, massive data archive systems, and high-speed network routers and switches. Each site also provides specialized systems for data analysis and visualization. While serving their respective user communities, NAS and NCCS are integrating processes such as requests for computing time and user account management: and building capabilities for large file transfer and cross-center data backup and recovery.

NASA Advanced Supercomputing Facility

The NAS facility at NASA Ames Research Center, known worldwide for its innovation and expertise in HEC, was built in 1986, following a long history of computing leadership at Ames extending back to the early 1950s. The NAS facility provides high-end computing resources and services for all mission directorates and the NASA Engineering and Safety Center, and serves time-critical Agency needs such as in-orbit Space Shuttle analysis.

Over the past two years, the NAS facility has undergone a major transformation; by taking advantage of vacant computer rooms at Ames, it has expanded from one computer floor to four, and nearly doubled supercomputing floor space, electrical power, and cooling. We have also transitioned to a multi-vendor environment, and expanded peak computational capacity from 62 to over 700 teraflops. In the process, both the primary and new computer rooms were retrofitted to handle the dramatically increased power consumption and cooling requirements. New 90- and 450-ton chillers were installed with associated pumps and plumbing, and the power complex was upgraded to include vast new wiring arrays and the largest power switch west of the Mississippi River.

Following these facility upgrades, NAS now operates four supercomputers: Pleiades, Columbia, RTJones, and Schirra, described in the *HEC Operations* section (page 5). In early 2008, after an extensive upgrade of the its visualization laboratory, NAS also installed the hyperwall-2 visualization and data analysis system, highlighted in the *Data Analysis and Visualization* section (page 13).

NAS will continue to enhance its high-end computing and visualization resources, supported by comprehensive user-focused services, to ensure high productivity and further mission success for HEC Program users and the Agency.

NASA Center for Computational Sciences Facility

The NCCS facility at NASA Goddard Space Flight Center was formed in 1990 with the arrival of its first Cray supercomputers. NCCS now supports modeling and analysis activities for Science Mission Directorate (SMD) users in Earth Science, Astrophysics, Heliophysics, and Planetary Science. Carrying on a role dating from the 1960s to provide computing and data services to the Agency's science community, the facility enables users to run the sophisticated models needed to make the best use of NASA satellite observations, prepare for upcoming flight missions, and support national and international scientific field campaigns. The NCCS facility operates a high-performance computing environment comprising hardware, software, networks, storage capabilities, and tools—and supports user consulting and training activities.

The primary NCCS supercomputer resource is the Discover cluster. Installed in 2006 and augmented several times since, this system has increased NCCS computing power five-fold and storage capacity nearly six-fold. Discover includes a multi-tiered storage system that supports data-intensive scientific computation through a large, online General Parallel File System, and a data migration facility for long-term storage and preservation of valuable project and user data. In addition, NCCS has implemented a Data Portal to distribute data to a broader set of scientists and engineers, and to enhance data sharing; and upgraded its analysis and visualization system to provide tools for model development and validation, and for performing science using model results.

NCCS continually refines and updates its data-centric HEC system architecture and service offerings to support the missions of advancing scientific research and understanding Earth, the solar system, and the universe.

HIGH-END COMPUTING SUPPORT SERVICES

NASA's HEC Program provides a full complement of services to support the Agency's scientists and engineers through the entire life cycle of their projects. Users require both capacity and capability computing, as well as batch and interactive service, including robust job scheduling and monitoring. Program support includes HEC operations, data management and archival storage, high-speed networking, 24x7 user support, application services, and data analysis and visualization—all targeted to help users make effective use of the HEC resources. Each facility optimizes its delivery of these services to best meet the specific needs of its user base.

HEC OPERATIONS

NASA's HEC operations services create the foundation for conducting large-scale modeling and simulation projects. These services provide not only massive computing power but also a wide array of tools ranging from compilers and debuggers, to job and workflow management. These services must address the challenges presented by increasingly complex models running on massively parallel architectures, and must help users make the most of both multi-core and specialty processors, such as graphics processing units. We meet these challenges through careful architecture planning and by implementing system hardware and software upgrades only after meticulous testing, to ensure maximum value with minimal disruption to users.

NAS HEC Operations

To keep pace with the growing needs of Agency users, NAS regularly evaluates new and emerging supercomputing architectures, and acquires and integrates those systems deemed to provide the best value to NASA. In early 2006, with the 10,240-processor Columbia supercomputer (62 teraflops, peak performance) at maximum utilization after about 18 months of operation, the NAS Technology Refresh (NTR) team began a four-phase process of enhancing NASA's HEC capacity. As part of the formal NTR evaluation process, we acquired a next-generation IBM POWER5+ system, named Schirra (320 dual-core processors), in spring 2007 to gain insight into the performance of this architecture on the NASA workload and to determine its feasibility to meet upcoming requirements. Later in 2007, an SGI Altix ICE system, named RTJones (1,024 quad-core Intel processors), was added to the NAS environment to support the Aeronautics Research Mission Directorate. In addition, Columbia was augmented to 14,336 processors and 89 teraflops (Tflops) peak performance—a 40% increase in capacity.

By spring 2008, NAS experts had gained extensive experience with all candidate HEC architectures, and determined that the SGI Altix ICE would provide the best overall value to the Agency. That fall, NAS and industry partners finished building the resulting Pleiades supercomputer, and integrating it with RTJones, establishing a system with 51,200 processor cores and 609 Tflops peak performance. At over eight times the capacity of Columbia's initial configuration, Pleiades exceeds the NTR goal to increase total sustained computing capacity at least four-fold every three years. Pleiades also achieved 487 Tflops on the HEC industry's LINPACK benchmark, making it the third most powerful supercomputer on the November 2008 TOP500 list. This ranking, combined with the November 2008 Green500 list ranking, made Pleiades the second-most powerful and energy-efficient general-purpose supercomputer in the world.

The Pleiades supercomputer, with a peak performance of 609 teraflops, is the third-fastest general-purpose supercomputer in the world (November 2008 TOP500 List).

Pleiades is already having a tangible impact on NASA's time-critical science and engineering challenges. For example, the Agency is saving thousands of hours and millions of dollars in experimental tests through computational analyses of Ares I Crew Launch Vehicle stage separation events and possible designs of the Orion Crew Exploration Vehicle thermal protection system.

NAS also refines its HEC environment for maximum reliability. For example, a unique weighting approach has been implemented to optimize data traffic movement to and from Pleiades through the nearly 20 miles of InfiniBand (IB) fabric that includes new fiber optic IB technology. Engineers customized the IB routing algorithm to reduce message contention and improve system transfer rates within a modified 10-dimensional hypercube. Due to the scale of Pleiades'

IB fabrics, a failover mechanism was implemented in MPT (SGI's version of Message Passing Interface), which can ride through temporary or permanent hardware failures between computer node end points. This also facilitates robust interconnections between Pleiades, Columbia, the mass storage systems, and the hyperwall-2 visualization system.

To ensure that users have access to the most effective resources, our experts also develop custom software tools and work with security experts on advanced methods to thwart regular attempts to break into NASA computers. NAS has developed and deployed a password system for the secure HEC enclave that enforces strict, uniform password rules to conform to NASA and federal regulations. The security team conducted a trade study and purchased a customizable, enterprise security management system for processing and correlating complex, interrelated, security-relevant information (such as logs and monitoring). This system feeds an extensible dashboard and sends alerts to flag potential intrusions.

NAS' near-term goal is to make Pleiades a stable production system by early 2009. In the longer term, NAS is preparing for a 10-petaflops HEC environment in 2012—a more than 10-fold increase over today's capability. Emerging innovative architectures and systems will be strategically leveraged to offer users the most effective supercomputing platforms and environments for NASA's computational challenges.

NCCS HEC Operations

NCCS provides computing resources for the Science Mission Directorate's (SMD) large-scale scientific and engineering models and simulations. NCCS computing services are uniquely configured to meet the data-intensive nature of these computations. SMD users require services to support development and execution of modeling and simulation codes.

During the last two years, NCCS has greatly increased its computing power, evolving to a more homogeneous, Linux-based cluster environment with vast, shared online disk storage. Compared with previous NCCS systems, this commodity approach offers greater computing and storage resources to the SMD user community, and reduces systems management complexity by implementing a single operating system. The current production platform, a Linux cluster named Discover, has 6,656 processor cores delivering 65 Tflops, with 10.8 terabytes of main memory and a 20-gigabit-per-second InfiniBand interconnect. This approach also allows greater flexibility in implementing incremental upgrades annually within the available budget. The resulting system enables large-scale ensemble runs, and storage and distribution of the data output. We also provide a test environment to evaluate system configuration changes and new user applications. Systems experts continually use benchmark information to determine required processor and I/O performance for model runs, and establish file organization approaches to optimize computational throughput.

The Discover cluster at NCCS. Discover's 6,656 processor cores yield a combined peak performance of 65 teraflops. During 2009, NCCS anticipates more than doubling Discover's capacity with an additional 8,192 processor cores, all based on Intel's new "Nehalem" processor.

Code development, the first stage in the life cycle of large-scale scientific applications (models), places great demands on scientists' time. NCCS provides a development environment where evolving models can be run. This environment includes a full complement of tools for developing, managing, debugging, and testing the code base of these applications. It also provides performance analysis tools and libraries, consulting services, and next-generation knowledge-based collaboration tools that enable developers to share their experiences. NCCS also maintains a code repository for SMD models, along with code versioning software to support modifications and upgrades. Access to the repository is granted to both NASA code development personnel and external collaborators. Our effective code development services support complex community-based model development and validation activities, and move codes more quickly from development to the production environment.

With the growing national interest in climate change research, contributions of NASA Earth Science programs to this field will continue to increase. NCCS has evaluated SMD's scientific and engineering initiatives and determined the HEC resource capability and capacity needed to meet those requirements. NCCS will continue to expand its Linux-based cluster(s) and enhance job-scheduling capabilities to address the increased processor count and diversity of job submissions. We are preparing to deliver unique and improved capabilities with increased computational and I/O performance to support complex model execution and data production, and meet the real-time execution needs of mission support operations. NCCS is also exploring enhanced workflow tools to allow seamless linkage between models, as well as use of "model recipes" and lessons learned from other scientists and engineers. We will support these workflows with global HEC resource management and automatic data migration on and off compute platforms. As time-to-solution for modeling and simulation remains a fundamental limitation, NCCS and their research partners will continue investigating opportunities to exploit multi-core parallelism on commodity processors and integration of powerful numerical accelerators.

DATA MANAGEMENT AND ARCHIVE

HEC data management and archive services address users' needs to store, retrieve, move, share, and preserve the vast amounts of invaluable scientific and engineering data produced by NASA and partner agencies. Data policies, including those for data aging and disposition, together with user needs for increased storage capacity to keep pace with computing power, guide planning for future data management and archive capabilities and capacity.

NAS Data Management and Archive

The NAS facility provides users with 16 petabytes (PB) of tape archive storage and 300 terabytes (TB) of online disk cache spread across three main archive systems. This is in addition to approximately 3 PB of disk storage attached to NAS supercomputing and data analysis systems. NAS' goal is to create the most robust and flexible high-performance storage environment possible.

For special projects, NAS storage experts create custom file systems to hold large amounts of temporary data, typically exceeding many terabytes. As an example, Science Mission Directorate (SMD) users working on ECCO2 global high-resolution ocean data syntheses required dozens of terabytes of disk storage to integrate their models and analyze results. We greatly expanded primary and secondary storage allocations and cross-mounted these to the archive storage system, substantially reducing bottlenecks during ECCO2 data generation and analysis. Custom training is also available to all NAS users requiring help with data management.

The Lou archive system housed at the NAS facility stores data generated on HEC Program supercomputers. In 2009, new archive systems will increase storage capacity seven-fold.

NAS has seen a steep growth in demands for computing time over the last several years—approximately 85% each year since 2004. This growth, in turn, places a heavy strain on the existing archive file systems. This trend is largely due to the increasing resolution of simulation runs and larger-capacity, higher-throughput supercomputers such as Pleiades and Columbia. In 2004, users transferred data into the archive at an average rate of 1.8 TB per day—today that number is more than 12 TB per day.

Over the last two years, our data management experts have made several key changes to improve long-term storage and transfer of data to archive systems. One improvement was to increase the size of the archive disk cache 10-fold by reusing older disk arrays from HEC platforms. Another improvement was to split the archive system, Lou, into three separate systems—one for SMD, one for the other three NASA mission directorates, and the third dedicated to testing and hyperwall-2 data analysis and visualization. This change helped streamline user access to data and improved transfer reliability. NAS also recently installed an InfiniBand (IB) network connection between Pleiades and the archive systems. As a result, users can now use batch job scripts to automate a file transfer from any computing node to the archive (via a file transfer protocol such as Secure Copy or bbftp). Previously, Pleiades users could only do a network file transfer from the front-end or bridge nodes.

In addition to hardware upgrades, both NAS and NCCS have worked to establish better redundancy and disaster recovery capabilities. NAS sends backups of system data to NCCS to ensure it is available in a second physical location in the event of a natural disaster or other unforeseen event. In turn, NCCS sends a large portion of their data to NAS. Each facility periodically stores encrypted files for one another to provide redundancy of key system information.

Several upgrades to the archive systems are in progress, including replacement of tape drives and storage silos no longer supported by the vendor. These upgrades will make more efficient use of space and take advantage of the latest tape densities and technologies. Following extensive testing, NAS has selected new hardware manufactured and supported by Spectra Logic, which will significantly reduce operating costs and floor space required—from approximately 800 square feet to only 70 on the main computer room floor, while increasing storage capacity seven-fold.

NCCS Data Management and Archive

NCCS addresses the full range of data management and stewardship requirements of the SMD user community. Data created and used by large-scale scientific codes and engineering simulations is accessible by all of the facility's HEC platforms and users, regardless of their physical location. This allows output from one computational run to be input to another without the user having to explicitly copy data from one system to another. Likewise, core observational and engineering datasets are available to SMD users, external researchers and collaborators, and the general public.

NCCS allows for sharing of data by users other than the data owners. This is especially beneficial when analyzing scientific data from emerging production codes, and when drawing from simulation and satellite observation data. To provide concurrent high-speed access to user files from every platform in the NCCS environment, a data store of 1.2 PB is visible to all systems via the Global Parallel File System (GPFS).

To facilitate large collaborative projects, we make it easy for users to set up project data spaces accessible by designated team members from various institutions, including public access for those outside NASA. Access to these data spaces is provided from within the Discover cluster, and distribution of project datasets are supported via the Data Portal. Currently, a copy of public data is stored to the local Data Portal within a local GPFS environment. With an increasing demand for sharing scientific results pertaining to climate change research, NCCS has implemented a method for the Data Portal to distribute data without needing to make local copies.

Increases in computational model complexity and number of job runs generate tremendous data management and archive requirements. NCCS combines heuristic evaluations of future requirements with actual storage usage statistics to develop capacity plans, and uses benchmark information to determine required I/O performance and establish file organization/striping approaches.

NCCS archives datasets using its Data Migration Facility (DMF) mass storage system, which currently holds up to 16.5 PB of data. Users define data requiring long-term storage, and DMF automatically migrates data from online disk cache to tape. Since 2006, NCCS has more than doubled the size of the DMF cache to better support increased archive demand commensurate with the increased computational capabilities. The addition of some 200 TB of online disk cache

The NCCS data archive infrastructure includes two linked StorageTek SL8500 silos. With an additional nine silos in other locations, NCCS has a total archive capacity of 16.5 petabytes.

to the DMF system, along with an additional 850 TB of online disk to the GPFS, allows users to maintain critical datasets online and tailor I/O performance to significantly reduce data retrievals from tape. Users manage and monitor archived data for longer-term need. The DMF allows deleting user datasets when no longer needed; transferring ownership when users leave projects; and proper disposal of data at the completion of projects.

The NCCS facility also has a unique stewardship requirement to provide long-term storage capacity to preserve invaluable scientific and engineering data produced by NASA and partner agencies. Project teams manage these datasets to provide continuity of responsibility, with formal agreements to ensure appropriate data retention.

Although limited search and query capabilities exist today, NCCS will incorporate additional data sharing and publishing capabilities in the future. These services will enhance data discovery, making it easier for scientists and engineers to locate relevant data to support projects. Implementation of these services will include enhanced metadata and supporting metadata search capabilities. Sharing and publishing rules and the related metadata will be managed and modified though a database, and will not require scientists and engineers to modify individual files.

Other upcoming plans include establishing capabilities and procedures for failover of critical applications from one facility to another.

HIGH-SPEED NETWORKING

The HEC Program's high-speed networking services provide fast, large-capacity connections to support exponentially increasing data transfers between computing resources and NASA's distributed user base. High-performance connectivity is supplied to and utilized by NASA centers and our university and research partners. Through in-depth analysis of traffic flows, the Program provides end-to-end (supercomputer-to-desktop) networking analysis and support for NASA and its partners. Networking experts develop solutions to meet the exponential growth in HEC network traffic, making multi-terabyte data transfers seamless for users.

NAS High-Speed Networking

Network engineers at the NAS facility, working together with NAS systems engineers, have for the past two years focused on developing innovative approaches for boosting network performance for the Agency's challenging applications and larger data transfers. This work involves the exploration of new technologies capable of maximizing network performance to and from NAS resources.

NAS network engineers have also been working closely with users, both remotely and through site visits, to help them use existing bandwidth more efficiently by: optimizing multiple aspects of end-to-end flows; tuning user systems; working with wide-area network (WAN) service providers such as National LambdaRail, the NASA Integrated Services Network, and Internet2; and working with user site infrastructure support teams to identify and remove bottlenecks along the network path. NAS-developed tools also help users take better advantage of existing HEC network capabilities. In one case, NAS engineers, in collaboration with engineers from NCCS, helped scientists working on the 3D Cloud-Resolving Model project (Goddard Cumulus Ensemble) increase data transfer rates by deploying NAS' version of the open-source file transfer application bbftp and making some end-system adjustments. This assistance resulted in a 54-fold improvement in network throughput performance and dramatic time-savings for the project.

Recently, our network engineers have tested 10-gigabit-per-second (Gbps) and faster firewall devices, which will help improve the security of HEC resources without creating bottlenecks in HEC traffic flows. Engineers are proactively looking for technologies to accelerate the supercomputer traffic flows to and from users. The latest development on this front is with the Cisco Wide Area Application Services, an appliance system capable of helping users achieve high performance with minimal adjustment to their desktop or local server system.

Network monitoring is another focus area for NAS. By using both custom-designed and off-the-shelf network analysis tools, engineers monitor network traffic conditions in real-time, and to look for historical trends. Recent innovations in the monitoring arena from cPacket and NetOptics include 10 Gbps "taps" that passively monitor network activity for trending, performance monitoring, and troubleshooting needs. The team is also investigating more effective ways to measure and evaluate increasingly large flows, using approaches that can readily scale above 10 Gbps while retaining high fidelity for smaller flows.

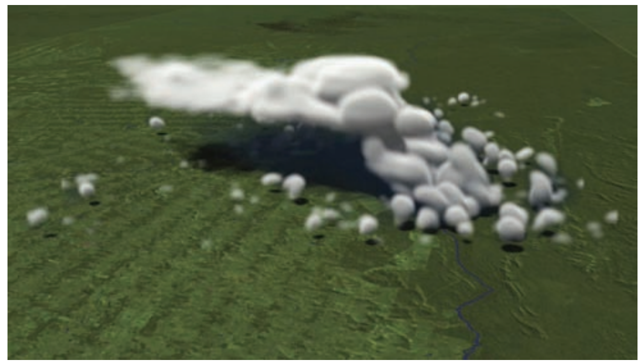

This visualization shows a three-dimensional, high-resolution simulation of a convective cloud system over South America. NASA Ames Research Center and NASA Goddard Space Flight Center network engineers helped scientists working on NASA's 3D Cloud-Resolving Model project dramatically increase their data transfer rate.

NAS is also focusing on enhanced network management. Using custom-designed and commercial off-the-shelf (COTS) network management tools, engineers manage a highly complex network environment. Examples of custom-designed tools include: the web-based Network Inventory Management System that allows more effective management of network resources; a web-based Access Control List (ACL) Management tool to more efficiently manage and troubleshoot ACLs; a Dynamic Host Configuration Protocol script that automates notifications to networking staff regarding users that exceed the allowed time on the system hardening network; and a packet size mismatch checking script to identify Maximum Transmission Unit errors. COTS tool examples include the ManageEngine network analysis tool suite, which has advanced network monitoring and management features, and Netflow Tracker to enable real-time monitoring of NetFlow data.

The NAS networking team recently completed three-year roadmaps for both WAN and local area network (LAN) environments. Required technology enhancements include lambda switching, disaster recovery, user-driven performance assessment tools, secure InfiniBand (IB) over WAN, and IB management. These advances in emergent network technologies are specifically targeted to supply the high-performance networking necessary to keep pace with the increasing computing capabilities. Further, tools such as these support effective management of complex networking environments.

By continuing collaborations with other research and development (R&D) partner networks, advanced networking vendors, and companies in Silicon Valley, our network engineers will remain on the forefront of emerging technologies to ensure that NASA's unique computational requirements are met. Important contributions include involvement in and representation within the Joint Engineering Team (a subcomponent of the interagency Networking and Information Technology Research and Development program), which coordinates networking activities, operations, and plans among multiple federal operational and research networks. NAS also reports on issues that affect the networking research community and makes recommendations to help mitigate those issues, especially related to security across government networks.

NCCS High-Speed Networking

NCCS networking efforts include both direct support of existing HEC user activities, and R&D in advanced communications technologies and protocols for future HEC production data flows. Over the past two years, several significant networking R&D achievements have come to fruition. A 10-Gbps coast-to-coast network was established between the University of California, San Diego and NCCS in Greenbelt, Maryland. This initial 10-Gbps capability is an important milestone toward establishing connectivity from NASA Goddard to the external high-performance network community. This effort earned a NASA group achievement award.

Also in 2006 at the annual Supercomputing Conference (SC06) in Tampa, Florida, the network team demonstrated use of the National Science Foundation's DRAGON Project's Xnet capabilities to dynamically stream uncompressed high-definition video from NASA Goddard to the exhibit floor. This was critical to showing the flexibility of new dynamic circuit switch technology to provide off-campus connections with immediate and temporary access to high-speed WANs. At SC07 in Reno, Nevada, this team supported real-time 3D imagery, further illustrating the ability of dynamic networks to properly synchronize multiple data streams required for 3D presentations.

Beyond these R&D activities, since 2006, NCCS has been providing direct support for existing user activities. To improve high-speed data transfer for NASA's HEC users, the NCCS Science and Engineering Network (SEN) infrastructure was upgraded to 10 Gbps. Networking teams also offer specialized support for individual Science Mission Directorate projects. The joint effort of the NAS and NCCS teams to help network users from Goddard's 3D Cloud Resolving Model project produced the 54-fold improvement in data transfer performance through enhanced file transfer tools and end-system tuning.

Near-term network support activities are focused on further improving network performance and data transfer rates for HEC users. We are working closely with the user community and the NASA Integrated Services Network to support the transition of intra-NASA HEC production flows. NCCS plans to further enhance the SEN infrastructure, and is researching opportunities to provide further performance gains for HEC projects.

Building on the existing 10-Gbps Ethernet WAN and 20-Gbps InfiniBand switching fabric, the NCCS network service team will continue evaluating new technologies to support NASA's HEC user community. Of particular interest are the assessments of emerging 40- to 100-Gbps communication technologies. Industry leaders are collaborating on a 40-Gbps Live Network trial. NASA Goddard will also be testing the use of a 100-Gbps Ethernet testbed that complements and interconnects with Internet2, ESnet, Infinera, Juniper, and Level3 networks. Additionally, NASA Goddard is sponsoring a Small Business Innovation Research/Small Business Technology Transfer opportunity for an n x 10-Gbps Offload network interface card for NASA, National LambdaRail, and grid computing to implement new technologies and protocols that support HEC science and engineering applications.

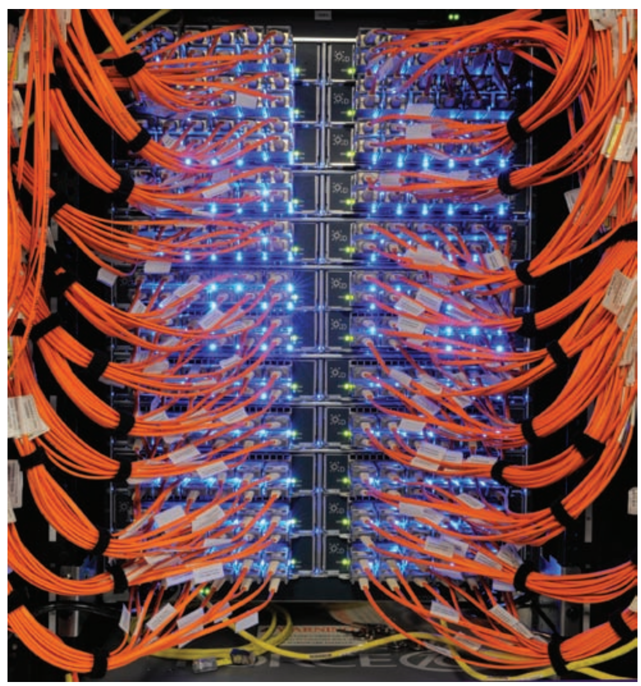

This tiered interconnect switch on the NCCS Discover Linux cluster manages data traffic over the 20 gigabit-per-second (Gbps) InfiniBand internal network. NCCS and the NAS facility maintain local-area networks and connect to each other via wide-area networks with 10-Gbps performance.

USER SUPPORT

User services staff provide direct assistance to the scientific and engineering user communities in every facet of their interaction with HEC resources—from setting up user accounts, to disseminating system information, to one-on-one problem solving and group training. Support ranges from resolving simple system usage issues to consulting on complex code optimization and advanced software techniques. The goal is to make effective use of the HEC resources and remove any obstacles in the way of user satisfaction.

NAS User Support

NAS provides tiered levels of support to all NASA mission directorate users: A 24x7 control room staff resolves basic system usage questions and coordinates end-to-end support (tier 1). Scientific consultants provide assistance in troubleshooting execution issues with user applications (tier 2), and applications experts perform significant code modifications using advanced software tools and techniques (tier 3).

In addition to providing tier 1 support, the control room team monitors the physical facility and HEC resources and manages computer jobs and queues—all to ensure a stable, productive computing environment for users. An important aspect of this team's work is the operational support provided for the aerothermal and debris analysis teams during Space Shuttle missions. Before and during each shuttle launch, NAS prepares and tests all HEC components to ensure engineers can provide computational analyses to mission managers to clear shuttles for landing. The ability to quickly reconfigure and reprioritize computing resources to assess potential shuttle launch damage is a key service to the shuttle program. The team also provides support for the hyperwall-2 visualization system.

NAS' scientific consulting services include application performance optimization; evaluating, installing, and customizing performance analysis software and tools; and specialized benchmarking of current and future HEC architectures to identify and leverage those best suited for the NAS computing environment. Application specialists examine performance characteristics of scientific and analysis codes and optimize them to enhance utilization of HEC resources and technologies. Small but high-impact code adjustments are provided routinely, and on request, NAS provides detailed analysis and advanced optimization services.

Over the last year, this work has greatly benefitted research applications for over 30 projects in all mission directorates. For example, a comprehensive optimization of USM3D—an important computational fluid dynamics (CFD) code used by the Aeronautics Research and Exploration Systems mission directorates for intensive aerodynamic analyses—improved the code's runtime two-fold, and reduced memory requirements by factor of 2 to support very large computational grids (over 100 million cells) as required for high-fidelity solutions.

HEC Program user services staff members are available 24x7 to answer basic system usage questions and coordinate end-to-end support.

Also within the last year, the applications team optimized two other important CFD codes, Phantom and OVERFLOW, which are heavily used by ESMD and SOMD for aerodynamic analyses of vehicle designs. In response to time-critical ESMD analysis needs, NAS applications experts reduced the run-times of these codes—saving nearly 2 million processor-hours and effectively freeing up more than an entire 512-processor node of the Columbia supercomputer. This work enabled completion of calculations that otherwise could not have been run.

Recently, NAS tool developers launched a web-based application for monitoring computing activity. Based on the more complex Heads Up Display (HUD) developed for NAS control room operators, this new "miniHUD" provides users an at-a-glance overview of the state of NAS supercomputers, enabling them to make informed decisions to manage their jobs. Users can "drill down" to get details on various nodes and subsystems, such as processor usage and job queue status.

With a new generation of supercomputers in place, NAS is now planning for a 10-petaflops HEC environment. To continue giving users world-class support as resources scale up, the user services staff is developing new and improved processes, tools, documentation methods, and training programs. We continue to work with other NASA centers on an integrated approach to handling user account creation and requests for computing time, and incorporating Agency initiatives (such as the NASA Account Management System and NASA Consolidated Active Directory) to ensure a more uniform and secure environment for users.

NCCS User Support

NCCS user services provides tiered levels of support to meet the full range of Science Mission Directorate (SMD) needs. Level 1 services, provided by NCCS and the Software Integration and Visualization Office (SIVO), include account support, help desk service, and access to system documentation along with system information services such as email notices, teleconferences, and forums. Level 2 and 3 services include consulting and training to help users make the most effective use of HEC resources. This includes supporting the use of modeling frameworks for greater portability and/or extensibility and improving code performance through optimization. NCCS User Services also coordinates Agency and center-provided services that are outside of NCCS but critical to providing a complete and responsive computing environment for advancing SMD projects.

NCCS and SIVO consultants provide expertise in computational science and scientific and engineering code development, including support in numerical techniques, software design and implementation, code parallelization using MPI or OpenMP, code porting, and performance optimization. Consulting services include both level 2 response to complex user queries and level 3 assistance to scientists and engineers requiring more extensive support.

NCCS training activity addresses all HEC services, platforms, software, and tools as well as a broad range of topics related to developing accurate and efficient scientific codes. Support is also provided for understanding and using data resources both within NCCS and at external data centers.

Over the last two years, significant user support efforts have been applied to evolving the system configuration to meet growing SMD computational needs. These needs include greater job scaling, increased memory, longer execution periods, and meeting product delivery timelines in support of operational projects and spaceflight opportunities; NCCS has determined the system requirements to support these needs. The user services team has also modified job submission and queuing policies to address the increased processor requirements of the models, the processor range for concurrent job execution, and job completion expectations.

Our user services experts also took the lead in providing a timely and successful migration of scientific codes from a shared memory system to a new distributed memory cluster environment. They contacted each user being impacted, provided training, assisted in code porting and/or redevelopment, and tracked the migration of applications in incremental steps. Additionally, this team proactively addressed user data needs as they migrated from SGI's clustered filesystem, CXFS, to the Global Parallel File System; and all user codes and data were successfully migrated.

In response to the Agency's evolving policies, NCCS has proactively worked with NASA Goddard personnel to clarify procedures for granting foreign nationals access to HEC Program supercomputers. These efforts have demystified and streamlined the procedural requirements and enabled the NCCS help desk to more easily establish user accounts. NCCS participates in NASA's security policy team and will continue to update its practices accordingly, and coordinate with the NAS facility to ensure consistency across the Program.

Another user services role is ensuring that users can adequately exploit future NCCS capabilities. Multi-core processors present a challenge for the entire high-end computing community, and NCCS is working to identify and evaluate performance enhancements these technologies could provide. Also, migration to newer versions of the Earth System Modeling Framework greatly improved memory management for complex codes, allowing greater scalability and better performance. As time-to-solution remains a fundamental research limitation, NCCS and SIVO are investigating using numerical accelerators for computationally intensive portions of NASA's scientific models. NCCS also supports activities to extend the use of NASA models and data to the broader scientific community.

Looking ahead, our user services staff will continue to reach out to the SMD user community to address their needs for high-end computing, data storage, visualization and analysis, and data sharing. NCCS will provide users with access to its knowledge-based trouble tracking system to give them more insight into problem resolution activities. Moving beyond static reporting will facilitate better user participation and expedited problem resolution.

Both the NCCS and NAS facilities will also be involved in forming a user board to represent the interests of HEC users. Following the example of an SMD Computational Modeling Capabilities Workshop in July 2008, workshops for other mission directorates are being planned to broadly understand and prepare for NASA's upcoming HEC needs.

DATA ANALYSIS AND VISUALIZATION

The HEC Program's data analysis and visualization services enable NASA scientists and engineers to find meaning in the vast amounts of data in their computational models and observational datasets. Analysis tools harness the power of computer processing to filter, sort, search, and compare datasets, as well as to apply advanced statistical and other types of algorithms—all with the goal of discovering useful information in datasets. In turn, visualization tools leverage the human brain's ability to identify interesting features and patterns in images and animations, allowing scientists and engineers to more deeply explore data and more clearly convey results to colleagues and the public.

NAS Data Analysis and Visualization

To help users understand and interpret their results, NAS visualization experts capture, process, and render enormous amounts of data to produce high-resolution images and videos. These experts also develop and adapt specialized visualization solutions for the Agency's unique science and engineering problems. Working closely with users, the team customizes and creates new tools and techniques to expose the intricate temporal and spatial details of computational models.

The NAS visualization team has developed special technologies for moving large datasets directly to graphics hardware as they are generated so that they can be displayed and analyzed in real time. A cornerstone of this capability is the hyperwall-2 visualization system installed in spring 2008. This powerful system provides a high-speed, fully interactive environment that enables users to visualize, analyze, and explore high-resolution results and pinpoint critical details in large, complex datasets. The hyperwall-2 is a matrix of functionally interconnected graphics workstations and displays coupled directly to the NAS facility's high-end computers via InfiniBand. The 128-screen, quarter-billion-pixel flat panel display system measures 23 feet wide by 10 feet high, giving users a supercomputer-scale environment to handle the very large datasets produced by high-end computers and observing instruments. Powered by 128 Nvidia 8800GTX graphics processing units (GPUs) and 1,024 Opteron processor cores, the hyperwall-2 has 74 teraflops of peak processing power and a storage capacity of 475 terabytes.

The hyperwall-2 is also integrated with a NAS-developed, state-of-the-art concurrent visualization framework, which enables real-time graphical processing and display of data while applications are running. This capability is critical to supporting huge datasets that are difficult to store, transfer, and view as a whole, and delivers results that are immediately available for analysis. Most importantly, concurrent visualization makes it feasible to render and store animations showing every simulation timestep, which allows users to see rapid processes in their models—often for the first time. With concurrent visualization users can also do on-the-fly identification of computational problems or parameter adjustments needed for applications. Live production data can be pulled from the supercomputer to hyperwall-2 without hindering code performance, allowing users to reap the benefits of visualization services without slowing turnaround time of their analyses.

The 128-screen hyperwall-2 visualization system enables users to view, analyze, and explore their high-resolution modeling and simulation results.

Other inventive visualization capabilities being explored at NAS include efficient GPU techniques, multivariate data visualization, and out-of-core techniques. Efficient GPU computation and rendering techniques enable both concurrent and rapid-iteration post-processing visualizations. Multivariate techniques include "linked derived spaces," which allows linking and selecting subsets of 2D scatterplots to facilitate visual tracking of key data points across all variables. Out-of-core data management techniques enable exploration of datasets that are too large to fit in memory.

Over the past couple of years, NAS experts have applied these capabilities to many applications, benefitting projects in all mission directorates. For example, detailed animations created from time-accurate, 3D simulations of ignition conditions within the flame trench of the Space Shuttle launch pad have provided NASA scientists with valuable insights into both repair criteria for damage caused during a shuttle liftoff, and potential modifications for the single-booster Ares I vehicle. Concurrent visualization methods have played a key role in identifying inherently unsteady flow structures in data-intensive simulations of V-22 Osprey rotors; and advanced GPU techniques have accelerated rendering of particle traces for massive simulations of convection on the Sun's surface.

Future advancements, such as streaming of visualizations to remote users and development of large GPU clusters, will expand the capacity and breadth of these valuable services. Visualization will continue to become more tightly integrated into the traditional environment of supercomputing, storage, and networks to offer an even more powerful tool to scientists and engineers.

NCCS Data Analysis and Visualization

NCCS provides a spectrum of data analysis and visualization services to help scientists access and manipulate large amounts of data produced by simulations, experiments, and satellite observations. For visualization, NCCS provides dedicated computing resources and access to large display systems. In addition, analysis services go beyond visualization to support commercial analysis software and tailored analytic tools developed by scientists. NCCS tools enable seamless movement of data to the analysis, visualization, and Data Portal environments, and data discovery tools help users assemble the required datasets for analysis. Each of these services is also supported by consulting and training for users.

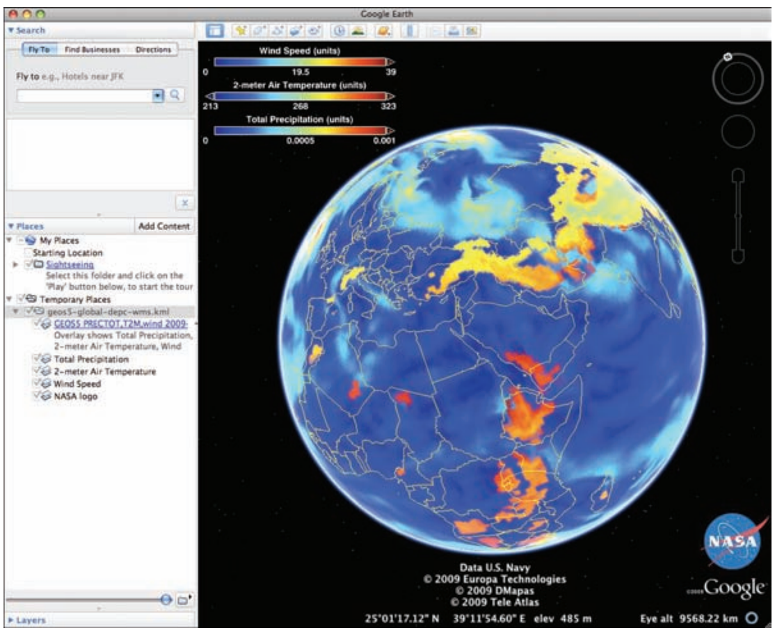

The NCCS Data Portal hosts the Web Map Service (WMS), which is provided by NASA Goddard's Software Integration and Visualization Office (SIVO). In this view, WMS visualizes multiple datasets from a run of the Goddard Earth Observing System Model, Version 5 (GEOS-5) and displays the results using the Google Earth interface.

Together with our scientific visualization partners, NCCS also provides specialized visualization support for Science Mission Directorate (SMD) projects. Several multi-panel hyperwall systems display data continuously at multiple locations around the NASA Goddard campus. These displays show current meteorological phenomena and highlight special features from SMD research. Multi-panel imagery allows scientists to visually assess before-and-after scenarios, and images can be transferred directly to user workstations for exploring the intricacies of their models and simulations. SMD users also employ NASA Goddard's Scientific Visualization Studio facilities to produce high-quality movies in support of education and public outreach. Complete with professional narration, these movies have appeared on many scientific television programs.

NCCS supports data analysis directly on the high-end computing platforms, and allows scientists to copy results of interest to their own servers or desktops. Using the NCCS Data Portal, collaborators can access data and perform limited analysis. While analysis on the supercomputers has proven extremely successful for model developers, and the Data Portal is a productive resource for many smaller applications, NCCS has also planned for capability enhancements to meet users' evolving scientific analysis challenges. Greater model complexity and resolution have increased data output volumes, making it more difficult for scientists to transfer data to their local systems. Analyzing these large datasets now requires more sophisticated, parallel analysis tools. Scientists often need to analyze data generated elsewhere in conjunction with external datasets, such as the NASA Goddard repositories of observing system data from many Agency satellites.

These growing requirements, combined with greater emphasis on collaborative research, led NCCS to install a new interactive, large-memory analysis platform with direct access to the entire NCCS global filesystem and data archive. This gives users fast access to large datasets needed for scientific analysis. Based on user interviews and knowledge of analytic techniques used at other supercomputing centers, NCCS has established a phased data analysis and visualization service development approach to incrementally address scientific user needs. This approach allows the flexibility for prioritizing capability enhancements with feedback from the user community. It allows scientists to continue with their current statistical model verifications using established tools such as Fortran, IDL, Matlab, or GrADS, while also forming the framework from which to explore new analytic paradigms for comprehensive scientific research.

FUTURE MISSION CHALLENGES

As an Agency-wide resource, NASA's HEC Program will support strategic directions set by the new presidential administration to: advance global climate change research and monitoring; mount a robust program of space exploration involving humans and robots; support the safe flight of the Space Shuttle to complete assembly of the International Space Station; and renew NASA's commitment to aeronautics research. The NAS facility serves all four of the Agency's mission directorates, while NCCS focuses on the Science Mission Directorate. Below are brief summaries of anticipated mission directorate plans for the near-term use of HEC Program resources and services.

Aeronautics Research Mission Directorate (ARMD)

As ARMD's largest consumer of HEC resources, the Fundamental Aeronautics Program (FAP) will continue to fully utilize its shares of the Pleiades, RTJones, and Columbia supercomputers at NAS. FAP uses these systems to enhance development of physics-based multidisciplinary design, analysis, and optimization tools for evaluating radically new vehicle designs and assessing the potential impact of innovative technologies on overall vehicle performance. High-end computing is enabling FAP's long-term, cutting-edge research addressing the concerns of modern air transportation. This research includes: improving aircraft performance while reducing noise and emissions; eliminating environmental and performance barriers to practical supersonic vehicles; improving mobility to meet greater demand for air transportation; and providing technologies to enable enhanced future space exploration capabilities.

Under ARMD's Airspace Systems Program, NASA's NextGen Airportal Project is harnessing the Pleiades supercomputer to develop technologies that will maximize single-airport capacity and improve the efficiency of multi-airport operations while maintaining or enhancing safety. Among the factors limiting runway capacity are the large spacing distances currently required on final approach to ensure that all aircraft avoid wake vortices from other aircraft. Researchers are developing a "fast-time" model that will accurately predict the precise location, movement, and decay rate of these wake vortices and provide datasets as simulation input for wake detection sensors. Research groups in the Aviation Safety Program are using the Columbia supercomputer to develop Advanced Satellite Aviation-Weather products for predicting in-flight icing; and to model polymer-based construction materials at the atomistic scale as part of the Aircraft Aging and Durability Project.

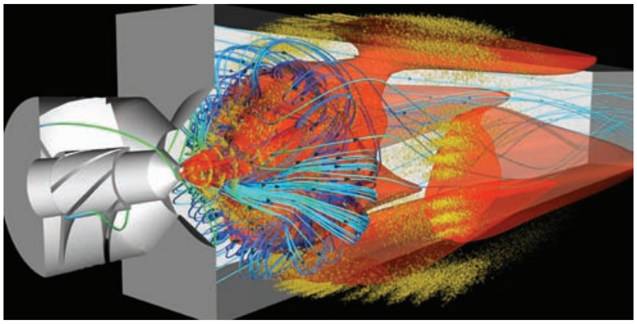

To help reduce harmful combustion emissions produced by aircraft, the Aeronautics Research Mission Directorate is assessing the effectiveness of software tools in predicting the presence of nitrogen oxide (NOx) and other emissions. The National Combustion Code was used to analyze the air flow through a lean-direct-injection combustor. Results show that the flow produces a recirculation zone that is critical to combustion stability but also produces NOx.

Exploration Systems Mission Directorate (ESMD)

In the coming years, HEC resources will continue assisting ESMD vehicle design, engineering, and mission planning. NASA supercomputers are supporting the Constellation Program's development of three next-generation space exploration vehicles: the Orion Crew Exploration Vehicle (CEV), the Ares I Crew Launch Vehicle (CLV), and the heavy-lift Ares V Cargo Launch Vehicle. As development ramps up, Constellation will need larger-scale simulations with finer grid resolutions and complex, time-accurate flow interactions to assess more intricate geometries and aerodynamic conditions.

Future support for CEV development will include aerothermal analyses for heatshield design, and assessments of design factors and abort conditions for launch abort system control. As the CLV design continues to be refined during later development stages, computational challenges will include aerodynamic analyses of: the vehicle's functional details such as fuel feed-lines, brackets, and umbilicals; guidance and attitude control system performance; and stage separation maneuvers with plume effects. Ares V aerodynamic analysis support will grow as the vehicle enters its major development phases. Computations will involve extensive aerodynamic database generation for each design cycle, as well as multi-species analyses of plume interaction and base heating effects for various engine configurations.

Additionally, the HEC Program has begun supporting ESMD's Lunar Precursor Robotic Program, which manages path-finding robotic missions to the Moon, leading the way

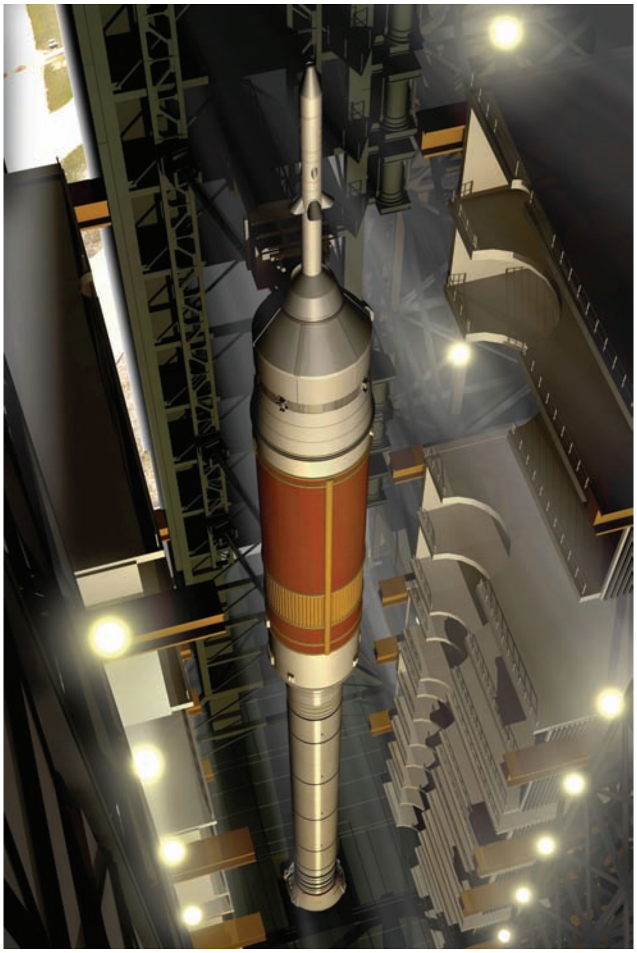

This artist's rendition shows the next-generation Ares I rocket being stacked in the Vehicle Assembly Building (VAB) at NASA's Kennedy Space Center in Florida. Simulations are helping to determine whether the VAB can safely manage combustion scenarios for the greater amounts of solid rocket booster fuel that Ares I and V will require.

to sustained human exploration of our solar system. The Lunar Crator Observatory and Sensing Satellite (LCROSS) mission—scheduled to launch with the Lunar Reconnaissance Orbiter in summer 2009—will determine whether water ice is present in a permanently shadowed crater at the Moon's south pole.

Science Mission Directorate (SMD)

Using Columbia and Pleiades at NAS and Discover at NCCS, SMD will analyze growing streams of data from new Earth-observing satellites and space missions. Across SMD's four divisions, models will need ever-higher spacial, temporal, and spectral resolution to match improvements in data resolution and to help NASA plan future observations.

Earth Science uses high-end computing to address problems in climate change and prediction. With greater resolution, NASA expects to better predict conditions such as hurricane intensity. Global climate simulations for the next Intergovernmental Panel on Climate Change assessment will include fully interacting atmospheric chemistry. Four-dimensional variational data assimilation is important for accurately incorporating precipitation observations into models and running Observing System Simulation Experiments for satellite design. Forecasting earthquakes and other solid Earth hazards requires daily simulations processing at least near-real-time, terabyte-sized data streams.

A high priority for Heliophysics is moving from today's 2D models towards 3D kinetic models of magnetic reconnection, a phenomenon that, among other effects, provides the energy release in solar eruptions. Forecasting space weather from initiation at the Sun to interaction with Earth and other solar system bodies means resolving each sub-domain and coupling domains across boundaries, with a minimum computational requirement equaling the sum of all domain calculations.

Planetary Science mission engineering employs high-fidelity space vehicle Entry, Descent, and Landing (EDL) simulations and landing site safety assessments. Expanding on techniques that have proven successful with the Phoenix Mars Lander, the Mars Science Laboratory (MSL) EDL simulations run four times as many particle trajectories at a higher frequency. Landing site safety analysis for MSL and a future Mars Sample Return mission entails hazard maps with 100 times the resolution used for Phoenix.

Astrophysics challenges include modeling the first stars and simulating how solar systems evolve from proto-planetary disks. Greater computing power is necessary for predicting gravitational waveforms from mergers of black holes with size differences of 20:1 to 100:1. Likewise, capturing star formation and other smaller-scale phenomena within cosmological simulations requires many more particles than is possible today.

New Earth-observing satellites and space missions being launched by NASA's Science Mission Directorate will generate more data at higher resolutions than ever before, presenting challenges for Earth and space science modelers. The Solar Dynamics Observatory shown above will help scientists better understand the Sun's influence on Earth and near-Earth space by studying the solar atmosphere on small scales of space and time and in many wavelengths simultaneously.

Space Operations Mission Directorate (SOMD)

With the approaching retirement of the Space Shuttle and a new generation of space vehicles on the way, the next few years are a pivotal time for SOMD. Throughout the remaining shuttle missions—currently planned through 2010—dedicated portions of NAS supercomputers and support staff must be on call from launch to landing to evaluate potential reentry risks posed by specific ice formation or foam debris damage sites, and to calculate heating on the thermal protection system to ensure a safe reentry. A mirrored Return to Flight data warehouse, co-located at NASA's Ames and Langley Research Centers, supports these on-the-fly calculations by rapidly transferring and disseminating hundreds of gigabytes of computational fluid dynamics data. SOMD will also rely on high-end computing resources for any future component redesigns or fuel tank processing improvements needed to keep the shuttle running safely and efficiently through its last missions. These efforts will become more important if the new administration opts to extend shuttle service beyond the planned retirement date.

Furthermore, high-end computing will contribute to the conversion of shuttle ground operations infrastructures for next-generation launch vehicles. For example, simulations are helping to determine whether the existing Vehicle Assembly Building at NASA Kennedy Space Center can safely manage potential combustion scenarios for the significantly greater amounts of solid rocket booster fuel that Ares I and V will require. Engineers are also performing analyses of shuttle and Ares ignition environments to determine whether modifications should be made to the current launch platform and flame trench configuration. Combining intensive, time-accurate plume and combustion simulations with large-scale models of ground operations infrastructures, such analyses present some of the most resource-demanding computations supported by NASA's High-End Computing Program.

This artist's rendition shows an Ares I rocket at Launch Pad 39B at NASA's Kennedy Space Center in Florida. Computational analyses of Space Shuttle and Ares ignition environments are helping NASA engineers to determine whether modifications to the current launch platform and flame trench configuration will be required to accommodate Ares.

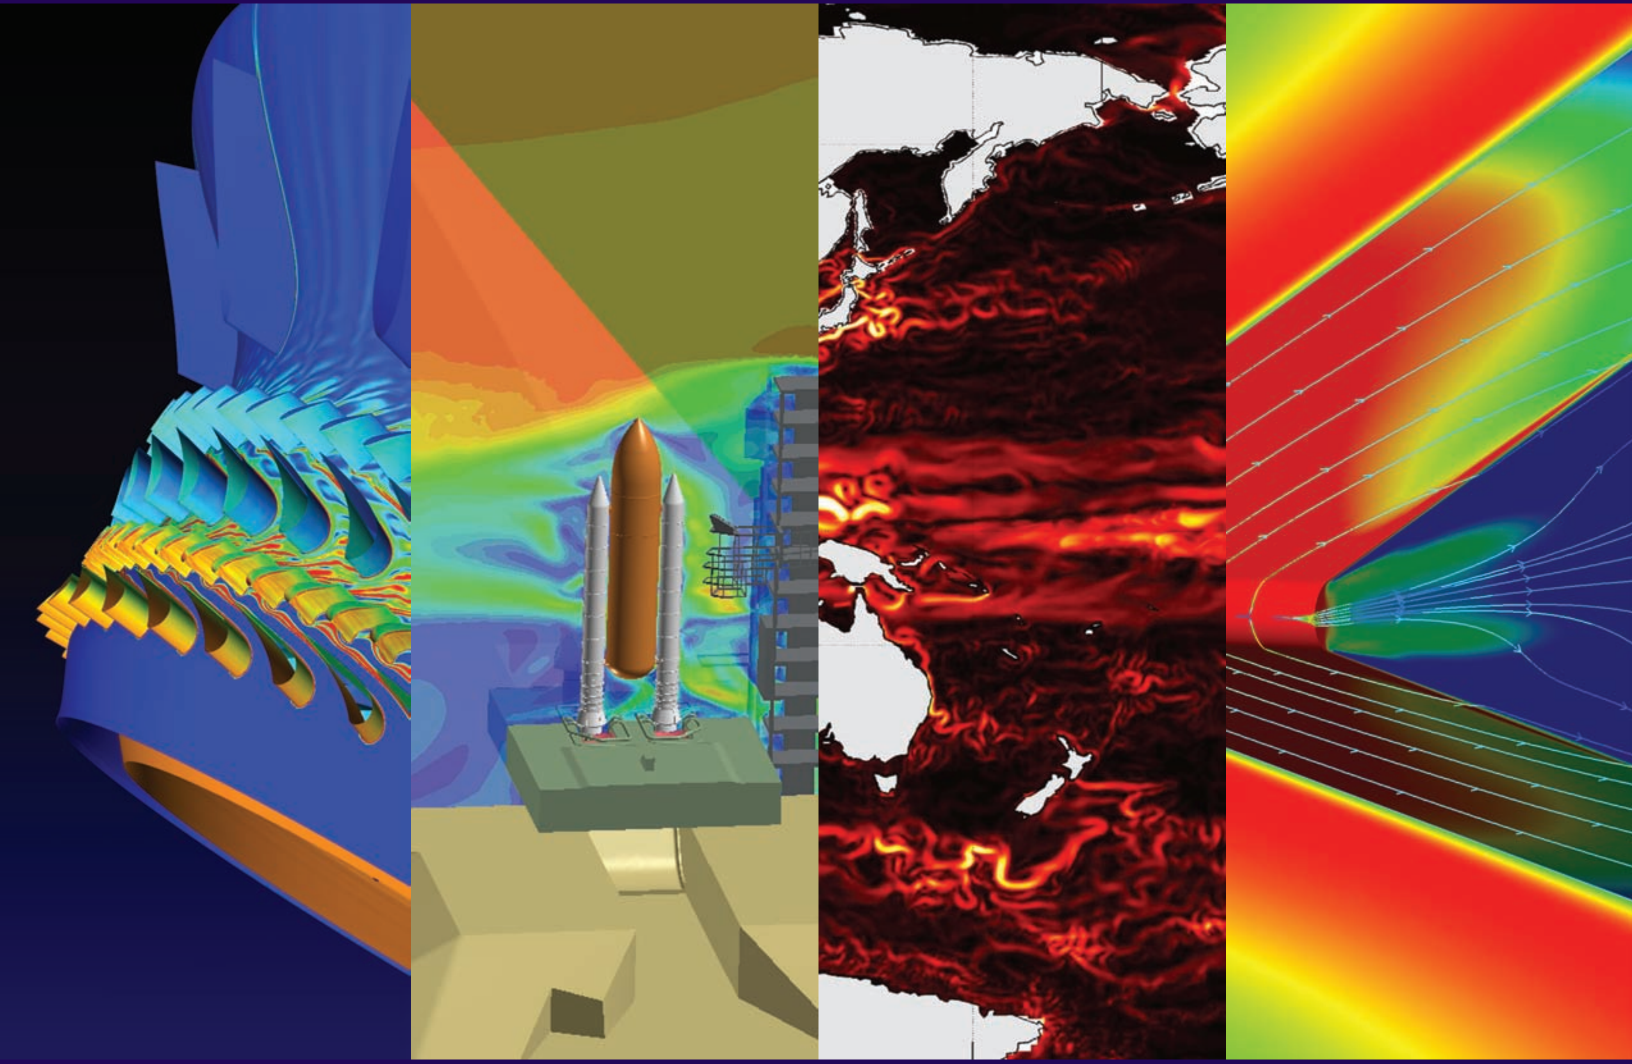

HIGHLIGHTS

SCIENCE AND ENGINEERING

This section presents 43 user projects from NASA's Aeronautics Research, Exploration Systems, Science, and Space Operations Mission Directorates, and the National Leadership Computing System Initiative, chosen because of their importance to the Agency, their impact during 2007 and 2008, and their technical maturity.